IN THE SPECIFICATION:

Please amend the paragraph beginning at page 6, line 8, as follows.

In section (3)(d), the development of a method for computing acoustic confusability between models is begun; in this section, confusability of densities is also discussed.~~.~~ This development proceeds by the following steps. In section (3)(e), the interaction between paths and densities is shown and discussed. In section (3)(f), an algorithm for computing confusability of hidden Markov models is disclosed. This method, as shown in section (3)(f), comprises the following: (i) constructing a product machine; (ii) defining the transition probabilities and synthetic likelihoods of the arcs and states of this product machine; and (iii) determining a probability flow matrix and relating the result of a certain matrix computation to acoustic confusability.

Please amend the paragraph beginning at page 13, line 26, as follows.

Thus, under these assumptions, the value $p(a(w) \mid h\ x)$ is now the quantity that needs to be determined, in either the continuous or discrete cases. This <u>disclosure</u> ~~discelosuer~~ will treat the continuous case in more detail later.

Please amend the paragraph beginning at page 14, line 17, as follows.

If there were only one single model $p(\cdot|x)$ for word x (that is, one single model for evaluating acoustic event probabilities, on the assumption that x was the word being pronounced), then $p(a(w) \mid x)$ would be the probability that this model assigns to the observation $a(w)$. But, in general, a given word x has many pronunciations. For instance, the word "the" may be pronounced "TH UH" or "TH IY," wherein the "TH," "UH," "TH," and "IY" are phonemes (also called "phones"). Phonemes are small, indivisible acoustic elements of speech; here UH represents a schwa sound, whereas IY represents a long e vowel. There are about 51 phonemes in English, although some speech recognition systems might have more or less phonemes. These pronunciations are referred to as lexemes or baseforms, and the following may be written:

$$x = \{l^1(x),\ l^2(x)\ ,\ \ldots\ ,\ l^{n_x}(x)\}, \tag{10}$$

to indicate that a word x admits of multiple pronunciations $l^1(x), l^2(x)$ and so on. Here n_x is the number of distinct pronunciations recognized for x and each $l^i(x)$ is one lexeme. Carrying this notation a little further, it is possible to write $l(x) \in x$ for an arbitrary lexeme $l(x)$ associated with the word x, and $\sum_{l(x)\in x}$ for a sum in which $l(x)$ varies over the lexeme set for x.

Please amend the paragraph beginning at page 60, line 10, as follows.

Turning now to FIG. 9, this figure shows a probability flow matrix ~~1000~~ that is populated using the product machine 830 of FIG. 8. Also shown in FIG. 9 is a column ~~1030~~ 930 that corresponds to the leftmost column of $((I - M) \mid I)$. Probability flow matrix ~~1000~~ contains probability flow matrix 600, which was shown FIG. 6. Additionally, the new state w_4 of the synthesizer model 810 of FIG. 8 causes entries 1001 through 1010 to be populated with probabilities. Determination of these types of probabilities has been previously discussed in reference to FIG. 6. From FIG. 9 and the previous discussion on Computational Caching, it can be seen that r_1 through r_9 will already be calculated when probability flow matrix 600 is used to determine acoustic confusability for synthesizer model 410 and evaluation model 420. Therefore, these may be held and reused when determining acoustic confusability from probability flow matrix ~~1000~~, which derives from synthesizer model 810 and evaluation model 420. This is a tremendous time savings, as r_{10} through r_{12} are the only values that need to be determined when probability flow matrix ~~1000~~ is used to determine acoustic confusability. For instance, it could be that synthesizer model 410 is the synthesizer model for "similar" and synthesizer model 810 is the synthesizer model for "similarity." The results r_1 through r_9 may be held and reused during the probability flow matrix calculations for "similarity." Likewise, the synthesizer model 810 could be the synthesizer model for "similar." The results for "similar" could be reused when computing acoustic confusability for "similarity." Note that the ordering of the states of the models will affect whether caching can be used for prefixes, suffixes or both.